



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation and Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein together with the Statement of inventorship and of right to grant of a Patent (Form 7/77), which was subsequently filed.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 8 July 2003

Patents Act 1977
Rule 16

THE PATENT OFFICE

10 APR 2003

The Patent Office

10APR03 E799277-1 D00611
P01/7700 0.00-0308262.5

Request for grant of a patent

NEWPORT

The Patent Office
Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

1.	Your reference	GB920030021GB1		
2.	Patent application number <i>(The Patent Office will fill in this part)</i>	0308262.5 10 APR 2003		
3.	Full name, address and postcode of the or of each applicant <i>(underline all surnames)</i>	INTERNATIONAL BUSINESS MACHINES CORPORATION Armonk New York 10504 United States of America		
Patents ADP number <i>(if you know it)</i>		519637001 <i>EE</i>		
If the applicant is a corporate body, give the country/state of its incorporation		State of New York United States of America		
4.	Title of the invention	RECOVERY FROM FAILURES WITHIN DATA PROCESSING SYSTEMS		
5.	Name of your agent <i>(if you have one)</i>	D P Litherland		
"Address for Service" in the United Kingdom to which all correspondance should be sent <i>(including the postcode)</i>		IBM United Kingdom Limited Intellectual Property Department Hursley Park Winchester Hampshire SO21 2JN		
Patents ADP number <i>(if you know it)</i>		6603898002 <i>EE</i>		
6.	If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and <i>(if you know it)</i> the or each application number	Country	Priority App No <i>(if you know it)</i>	Date of filing <i>(day/month/year)</i>
7.	If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application	No of earlier application		Date of filing <i>(day/month/year)</i>

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:
 a) any applicant named in part 3 is not an inventor, or
 b) there is an inventor who is not named as an applicant, or
 c) any named applicant is a corporate body.)

Yes

9. Enter the number of sheets for any of the following items you are filing with this form.
 Do not count copies of the same document

Continuation sheets of this form

Description 21

Claim(s) 6

Abstract 1

Drawing(s) 5

xs fm

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents
 (please specify)

11.

I/We request the grant of a patent on the basis of this application

D P Litherland
 Signature9 April 2003
 Date

12. Name and daytime telephone number of person to contact in the United Kingdom D P Litherland
 01962 816303

Patents Act 1977

Rule 15

THE PATENT OFFICE
D
16 APR 2003
NEWPORT

The
Patent
Office

16APR03 B129499-1 000611
P07/7700 0-00-0308262-5

Statement of inventorship and of right to grant of a patent

The Patent Office
Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

1. Your reference GB920030021GB1
2. Patent application number 0308262.5
(If you know it)
3. Full name of the or of each applicant INTERNATIONAL BUSINESS MACHINES CORPORATION
4. Title of invention RECOVERY FROM FAILURES WITHIN DATA PROCESSING SYSTEMS
5. State how the applicant(s) derived the right from the inventor(s) to be granted a patent By employment and by agreement
6. How many, if any, additional Patents Forms 7/77 are attached to this form? 2
7. I/We believe that the person(s) named over the page (and on any extra copies of this form) is/are the inventor(s) of the invention which the above patent application relates to.

D P Litherland
Signature

14 April 2003
Date

8. Name and daytime telephone number of person to contact in the United Kingdom D P Litherland
Tel: 01962 816303

Enter the full names, addresses and postcodes of the inventors in the boxes and underline the surnames

Robert Frank BUXTON
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

6916786002

Patents ADP number (*if known*)

David James FISHER
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

8359390001

Patents ADP number (*if known*)

If there are more than three inventors, please write their names and addresses on the back of another Patents Form 7/77 and attach it to this form

Stephen James HOBSON
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

8359408001

REMINDER

Have you signed the form?

Patents ADP number (*if known*)

Enter the full names, addresses and postcodes of the inventors in the boxes and underline the surnames

Paul HOPEWELL
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

83594⁴⁰001

Patents ADP number (*if known*)

If there are more than three inventors, please write their names and addresses on the back of another Patents Form 7/77 and attach it to this form

Robert Daniel MILLAR
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

8359424001

REMINDER

Have you signed the form?

Patents ADP number (*if known*)

Enter the full names, addresses and postcodes of the inventors in the boxes and underline the surnames

Peter SIDDALL
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

8359432001

Patents ADP number (*if known*)

Stephen Richard WALKER
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

8359457001

Patents ADP number (*if known*)

If there are more than three inventors, please write their names and addresses on the back of another Patents Form 7/77 and attach it to this form

REMINDER

Have you signed the form?

Patents ADP number (*if known*)

RECOVERY FROM FAILURES WITHIN DATA PROCESSING SYSTEMS

Field of Invention

5 The present invention relates to recovery from failures in data processing systems, and in particular to recovery components and methods implemented within computer programs and data processing systems.

Background

10 Even very reliable data processing systems can be susceptible to storage failures, such as disk failures and malfunctions or software malfunctions, that result in loss or corruption of data in primary storage. To avoid such failures resulting in permanent loss of data, it is known to 15 provide recovery capabilities including making backup copies of stored data and taking log records describing the updates to the stored data since the latest backup.

A number of communication manager software products, including IBM 20 Corporation's MQSeries™ and WebSphere™ MQ family of messaging products, provide facilities for storing messages in a data repository such as a message queue or database table during transfer of messages between a sender and a receiver. As with other data processing systems and computer programs, there is a need for solutions for recovering from potential 25 system or program failures to avoid loss of critical messages and to ensure that application program tasks can complete successfully.

In a message queuing system in which queue manager programs handle 30 the transfer of messages between queues, it is known for recovery facilities within the queue manager programs to recover a queue and its message contents when the primary storage used to hold its messages fails. The recovery facilities restore messages to the queue so that the final state of the queue is the same as at the time of the storage failure. These 35 recovery facilities recreate a message queue and a snapshot of its contents from a back-up copy of the queue, and then refer to the queue manager's log records to reapply changes to the queue. In such known solutions, queue managers must complete the recovery processing before any messages are retrieved from the queue, and before any new messages are added to the queue. This ensures that the state of the queue after recovery is the same 40 as the state of the queue at the time of the failure, and that message sequencing is not lost as a result of the failure.

However, a remaining problem with such solutions is the unavailability of the messaging functions and the message repository while

the recovery processing is in progress. Many applications require optimum message availability but have competing requirements for the messaging system to provide assured once-only message delivery. If an application is allowed to access a queue during the recovery processing, there is a danger 5 that a single message may be processed twice by the application. A bank customer who has funds debited from his account twice in response to a single funds transfer instruction would be very dissatisfied.

US Patent No. 6,377,959 issued on 23 April 2002 to Carlson describes 10 a transaction processing system that continues to process incoming transactions during the failure and recovery of either one of two duplicate databases. One of the two duplicates is assigned "active" status, and the other is maintained with "redundant" status. All incoming queries are sent only to the active database and all incoming updates are sent to both the 15 active and redundant databases. When one database fails, the other is assigned active status (if not already active) and continues to process incoming queries and updates during repair and restart of the failed database. Repair and restart of the failed database involves use of interleaved copy and update operations in a single pass through the active 20 database. The interleaving of incoming updates and copy operations is performed according to a queue thresholding method, which controls copy operations in response to the number of incoming transactional updates. The transaction processing system remains operational both during the failure 25 and recovery activities. Since a full replica is maintained, log records are only written when one of the databases fails, and access is not required to the failed database while that database is under repair. Although continuous availability is highly desirable, this solution has the significant processing and storage overhead of maintaining two complete 30 database replicas with interchangeability of the operating status (active or redundant) of each of the two database systems. Furthermore, replication generally does not protect against software corruption, and so recovery operations will be required in addition to replication in some circumstances.

35 US Patent Application Publication No. 2002/0049776 (published on 25 April 2002 for Aronoff et al) also relates to replicated databases for high availability. The document describes a method for resynchronization of source and target databases following a failure by restarting replication after recovery of the target database and purging stale transactions that 40 have already been applied to the target database during recovery.

An alternative approach is described in US Patent No. 6,353,834 issued on 5 March 2002 to Wong et al, in which a message queueing system stores messages and state information about the messages, clustered

together in a single file on a single disk. This system is intended to achieve efficient writing of data by avoiding writing updates to three different disks (a data disk, an index structure disk and a log disk). A Queue Entry Map Table is used to enter control information, message blocks and log records. US 6,353,834 refers to the use of existing RAID technology and duplicate writing of data, without which the described system provides no protection against storage failures which result in loss of the data held on the single disk.

10 Summary

Aspects of the present invention provide methods, data processing systems, recovery components and computer programs for recovering from failures affecting data repositories.

15 A first aspect of the invention provides a method for recovery from failures affecting a primary copy of a data repository, for use in a data processing system in which updates applied to the data repository during normal forward processing are applied within transactional units of work.
20 The method includes storing a secondary copy of data representing data items held within the data repository and updates applied to the data repository within transactional units of work. In response to a failure affecting a primary copy of the data repository, the secondary copy is used to identify a set of operations required for restoring data items and
25 applied updates to a primary copy of the data repository. The set of operations are analyzed to determine the state, at the time of the failure, of each unit of work corresponding to one or more operations of the identified set of restore operations. Restore operations of the identified set are then performed if performance is consistent with the determined
30 state of the corresponding unit of work, but restore operations for which performance is inconsistent with the determined state of the corresponding unit of work are discarded without being performed.

The above-described method enables more efficient recovery processing than methods which merely re-apply all updates in the sequence in which they appear in the log, while also maintaining transactional integrity.

A further aspect of the present invention provides a data communication system including: data storage for storing a primary copy of a data repository; secondary data storage for storing a secondary copy of data representing the data repository which secondary data is sufficient to recreate the primary copy of the data repository and data held thereon; and a recovery component for controlling the operation of the data communication system to recover from a storage failure affecting the

primary copy of the data repository. The recovery component is operable to control the data communication system to perform the method steps described above.

5 Methods according to the invention preferably include the step of saving to a cache a subset of the secondary copy of data. This subset corresponds to the identified set of operations required for restoring data items and applied updates. Subsequent to the step of determining the state of each unit of work, restore operations are retrieved from the cache and
10 applied to the primary copy of the data repository.

Preferably, restore operations for which the corresponding unit of work is determined to be neither committed nor in-doubt are deleted from the cache prior to applying restore operations. This ensures that restore
15 operations for which performance is inconsistent with the determined state of the corresponding unit of work are disregarded when performing restore operations.

The performance of restore operations preferably comprises:
20 performing restore operations for which the corresponding unit of work is determined to be committed; and performing restore operations for which the corresponding unit of work is determined to be in-doubt, and marking the data item to indicate that the unit of work is in-doubt.

25 In a preferred embodiment, the method includes deleting from the cache any pairs of updates within the set of restore operations, which pair of updates correspond to addition of a data item and retrieval of the same data item and which pair of updates was completed prior to the failure. This ensures that such pairs of updates are disregarded when performing
30 restore operations - avoiding unnecessary processing.

In a messaging embodiment of the invention, if a pair of updates to a message repository correspond to addition of a message and retrieval of the same message, and the pair of updates was completed prior to the failure,
35 the pair of operations can be performed together within recovery processing without risk of leaving the repository in an inconsistent state. Such 'add-retrieve' pairs of operations are identified when log records are replayed. The pairs of operations are either omitted from the restore processing (i.e. deemed to have been performed as a pair, since their
40 effects on the queue cancel each other out) or the pairs of operations are performed and committed outside of the scope of the Recovery Unit of Work. Each of these options avoids unnecessary processing and reduces the potential build-up of messages.

The above method mitigates a problem which affects many known communication solutions - which is the tendency for data to build up in repositories while recovery processing is being carried out. This problem can result in the repository (or structures within the repository) reaching 5 a 'full' condition. The results could be that some data communications are returned to the sender or build up at an intermediate network location, unless significant additional processing is carried out to prevent this.

According to a preferred embodiment of the invention, updates to a 10 message repository during normal forward processing of a messaging system include message send operations which add messages to the repository, and message retrieve operations which delete the messages. The 'message repository' in this context may be a message queue, a database table, or any other data structure which holds messages or message queues. Following 15 a failure which affects the message repository, send and retrieve operations are reapplied to the repository, by referring to a backup copy of the repository and log records. The log is read to identify operations required to restore the message repository, but these operations are deferred until a determination can be made of the state of each unit of 20 work corresponding to the identified operations.

Preferred embodiments of the invention enable recovery from primary storage failures in a shared-queue messaging system, including recovery of old messages (messages from before queue failure) onto shared queues from 25 backup copies of the queue and log records. The shared queues may be in use by one or more application programs processing new messages (messages sent to the queue after the failure) while old message repository updates are being restored from log records. This message recovery can be performed while also providing assured once-only delivery of messages by handling the 30 entire restore processing as a single unit of work.

Methods and recovery components as described above may be implemented within a computer program for controlling the performance of a data processing apparatus on which the program code executes. The program code 35 may be made commercially available as a program product comprising program code recorded on a recording medium, or may be made available for download via a network such as the Internet.

Brief Description of Drawings

40

Embodiments of the invention are described in detail below, by way of example, with reference to the accompanying drawings in which:

Figure 1 shows a message communication network, in which messages are transferred between queues on route to target application programs.

5 Figure 2 is a representation of a set of queue managers having shared access to a queue within a coupling facility list structure;

Figure 3 shows a sequence of steps of a recovery method according to an embodiment of the invention; and

10 Figure 4 shows a sequence of steps of a recovery unit of work according to an embodiment of the invention.

Detailed Description

15 A first embodiment of the invention is described below in the context of asynchronous message communication systems in which messages are queued in message repositories between the steps of a sender program sending the message and a retriever program retrieving the message. A failure of primary storage can cause loss or corruption of message data unless 20 recovery features are available to recreate the queue and to recover messages onto the queue. While applicable to other data repositories, the invention is particularly applicable to messages queues because such queues typically contain discrete independent items (the messages) which are added and then deleted, rather than the message being added, its content updated, 25 and then finally deleted.

As will be clear to persons skilled in the art, certain embodiments 30 of the invention are equally applicable in a database environment in which a failure can result in loss or corruption of data within a database table, and thus necessitate recreation of the database table and restoring of data 35 items into the table. Embodiments of the invention are also applicable in other data processing environments in which hardware or software failures necessitate recovery of a data repository, for example from backup storage and log records, and in which there is a need to minimize the loss of availability of the data repository while recovery processing is carried out.

Loss or corruption of data on a primary storage medium may result 40 from a hardware failure or malfunction, a software malfunction, or even a human error (such as an accidental deletion of a queue and all of its messages). For ease of reference, all of these different types of failure which affect a data repository will be referred to as 'storage failures' hereafter. The loss or corruption may affect only a single queue, or database table, or file, or the failure may affect more than one queue (or

table etc) such as multiple queues held within a single Coupling Facility list structure (see the explanation of CF list structures below). In typical cases, a failure affecting a CF list structure will affect all queues on the CF list structure rather than a single queue.

5

Messaging Environment

IBM Corporation's MQSeries™ and WebSphere™ MQ family of messaging products are examples of known products which use message queuing to support interoperation between application programs, which may be running on different systems in a distributed heterogeneous environment.

Message queuing and commercially available message queuing products are described in B.Blakeley, H.Harris & R.Lewis, "Messaging and Queuing Using the MQI", McGraw-Hill, 1994, and in the following publications which are available from IBM Corporation: "An Introduction to Messaging and Queuing" (IBM Document number GC33-0805-00) and "MQSeries - Message Queue Interface Technical Reference" (IBM Document number SC33-0850-01). The network via which the computers communicate using message queuing may be the Internet, an intranet, or any computer network. MQSeries and WebSphere are trademarks of IBM Corporation.

As is well known in transaction processing systems, a 'unit of work' is a set of processing operations that must be successfully performed together, or all backed out in the event of inability to complete the full set of operations, to ensure that data integrity is not lost. All operations within a unit of work are kept inaccessible from other processes, which may rely on the updates, until resolution of the entire unit of work allows all of the updates to be committed (all finalized and made accessible).

IBM Corporation's MQSeries and WebSphere MQ messaging products provide transactional messaging support, synchronising messages within logical units of work in accordance with a messaging protocol which gives assured once-only message delivery even in the event of system or communications failures. This assured delivery is achieved by not finally deleting a message from storage on a sender system until the message is confirmed as safely stored by a receiver system, and by use of sophisticated recovery facilities. Prior to commitment of transfer of the message upon confirmation of successful storage, both the deletion of the message from storage at the sender system and insertion into storage at the receiver system are flagged as uncommitted (in flight or in doubt operations) and can be backed out atomically in the event of a failure. This message transmission protocol and the associated transactional

concepts and recovery facilities are described in International Patent Application Publication No. WO 95/10805 and US Patent No. 5,465,328.

The inter-program communication facilities of IBM's MQSeries and
5 WebSphere MQ products enable each application program to send messages to
the input queue of any other target application program, and each target
application can asynchronously take these messages from its input queue for
processing. This achieves delivery of messages between application programs
that may be spread across a distributed heterogeneous computer network,
10 without requiring a dedicated logical end-to-end connection between the
application programs

Recent versions of IBM Corporation's MQSeries for OS/390 queue
manager software provide support for shared queues using OS/390 coupling
15 facility (CF) list structures as the primary storage for shared queues.
Messages on shared queues are stored as list entries in CF list structures.
Applications running on multiple queue managers in the same queue sharing
group anywhere in a parallel sysplex can then access these shared-queue
messages, with messages being accessed in the order of allocated primary
20 keys. From the viewpoint of the Coupling Facility, the allocation of the
primary keys is arbitrarily decided and associated with each message by the
queue manager. The queue manager sets the key for each message so that the
overall order is the correct order for retrieval (applying FIFO ordering
with exceptions, as described below).
25

Such shared access to specific queues has the benefits of high
availability through redundancy (tolerance to failures affecting one or
more queue managers within the group) and automatic workload balancing
since messages are retrieved by the next available application. This
30 provides a highly scalable architecture suitable for high message
throughput.

The present embodiment is applicable to the system architecture
described above - and indeed is beneficial since many applications running
35 in this environment require high availability - but embodiments of the
invention are also applicable where alternative storage structures are
used. Hereafter, the term message repository is used to refer to message
queues and other data structures in which messages can be held, whether
implemented in CF list structures, database tables or other known
40 structures.

As noted above, message queuing systems in the OS/390 operating
system environment provide support for shared queues that can be made
available to a queue-sharing group of queue managers via CF list

structures. System components, data structures and methods applicable to such systems, including a number of recovery features which are suitable for use within such systems, are described in the specifications of the following co-pending and commonly-assigned patent applications, each of 5 which is incorporated herein by reference:

- US Patent Application No. 09/605589 (corresponding to UK Patent Application No. 0009989.5 - Attorney reference GB920000031),
- US Patent Application No. 09/912279 (Attorney reference GB920000032),
- 10 • US Patent application No. 10/228615 (corresponding to UK Patent Application No. 0207969.7 - Attorney reference GB920010101),
- US Patent Application No. 10/228636 (corresponding to UK Patent Application No. 0207967.1 - Attorney reference GB920020001) and
- US Patent Application No. 10/256093 (corresponding to UK Patent Application No. 0208143.8 - Attorney reference GB920020015).

15 The embodiment of the present invention described below is compatible with the recovery features described in the above-listed incorporated references.

20 Methods and apparatus for implementing message queues within list structures and processing list structures, as well as solutions for differentiating between operational states using distinctive keys, are described in the specifications of the following co-pending, 25 commonly-assigned patent applications, each of which is incorporated herein by reference: US Patent Application No. 09/677,339, filed 2 October 2000, entitled "Method and Apparatus for Processing a List Structure" (Attorney reference POU920000043); and US Patent Application No. 09/677,341, filed 2 October 2000, entitled "Method and Apparatus for Implementing a Shared 30 Message Queue Using a List Structure" (Attorney reference POU920000042).

Figure 1 shows, schematically, a messaging network 10 in which messages are transferred between queues 20 under the control of queue manager programs 30 in a distributed network of computers 80. Sender 35 application programs 40 put messages to their local queue, and target application programs 50 retrieve messages from their input queue, and all of the work of transferring the message across the network to the input queue of the target application program without loss of persistent messages is handled by the queue managers 30. Each queue manager maintains a backup 40 copy 60 of its local queues and writes log records 70 to reflect updates whenever messages are added or deleted or their state is changed.

Figure 2 shows a group of queue managers 30 which have shared access to queues 100 held in a Coupling Facility (CF) list structure 110. The CF list structures are used to queue messages in both directions - to and from the queue-sharing group. In addition to the primary copy of the shared 5 queue, a secondary backup copy 60 is held on a disk 120. Backup copies of the queue, comprising queue definition information and information relating to all the messages held on the queue at the time of the backup, are saved periodically to the disk. Log records 70 are written to the disk 120 for each update to a queue within the CF list structure. The combination of a 10 backup copy and log records reflecting all updates since the last backup enables recreation of the primary copy of the queue in response to a media failure.

The log records contain an indication of the operation performed 15 (insert, delete, or update state), and the unique key for the relevant message which key is generated at the time the message is added to the CF. For insert operations (and for update operations in some implementations) the log record also contains the complete content of the message. Log records for delete operations do not contain the content of the database 20 records. In some implementations, only the information required to track changes is logged for update operations.

Recovery with improved availability

Some computer systems and applications can tolerate "out of sequence" 25 updates to data repositories. That is, the systems work correctly even if the sequence of updates in the repository does not accurately reflect the sequence in which the updates were added. This is true of some systems and applications, which use message queue managers to transfer messages to and 30 from queues when handling message delivery between application programs.

The inventors of the present invention have recognized that such systems and applications could benefit from improved availability by enabling new messages to be added to and retrieved from queues prior to 35 completion of recovery of the data on the queues following a failure. However, before this can be achieved, a number of problems must be overcome.

If an application is enabled to access a newly created queue in 40 parallel with old messages being restored to the queue by replay of log records, there is a danger that the same message may be processed twice by the application. For example, a message may be added to a queue, the addition operation committed, and then the message retrieved from the queue. In most cases, the message is deleted from the queue when the

retrieval operation is committed. If a queue storage failure then occurs, the queue can be recreated from backup storage followed by reapplying updates to the queue from log records. During log replay, the message is restored to the queue and becomes available to retriever applications when
5 the commit of the addition operation is replayed, and then disappears when the message retrieval operation is replayed. However, if application programs are able to access the queue during recovery, an application program may retrieve the message as soon as it becomes available (i.e. before replay of the message retrieval log record) and process a message
10 which has already been processed before.

The above sequence of events, and other examples, can result in unacceptable deviation from assured once-only message delivery.

15 A solution to this problem is described below, which can recover from a primary storage failure by recovering messages to shared queues while the shared queues are in use by an application which is processing new messages, without deviating from assured once-only delivery of messages. 'New messages' in this context are messages added to the queue for the
20 first time after a failure. 'Old messages' are those that were added to the queue prior to the failure and which are restored to the queue following the failure.

Recovery Processing within Recovery Unit of Work

25 In the present embodiment, the restore process is performed as a Recovery Unit of Work. That is, the sequence of steps of restoring messages to a queue and updating the state of messages on the queue from backup storage and by replaying the log are performed and committed within the
30 scope of a newly-defined unit of work.

For example, the actions of replaying an out-of-syncpoint message 'Put' operation (adding a message to a queue) or 'Get' operation (retrieving a message from the queue), or replaying commit of an
35 in-syncpoint Put or Get, are performed as in-syncpoint Puts and Gets within the Recovery Unit of Work. The Recovery Unit of Work covers the entire process of restoring messages to the queue and replaying operations which change the state of those messages.

40 A unit of work is a set of operations which must be performed together (or not at all) if the data affected by the set of operations is to be left in a consistent state at the end of performing the set of operations. A syncpoint is an identifiable point within processing at which data is in a consistent state, and syncpoints are recorded at the end

of each unit of work to record this point of consistency. Reference to recorded syncpoints enables a determination to be made of how far back in time to rollback processing in order to return to a point of data consistency. A single transaction can include a number of Put_Message and 5 Get_Message operations which are processed as a single unit of work. When the transaction is committed, all of the Put and Get operations within the unit of work are finalized such that messages Put onto a queue appear on the queue as retrievable messages and messages for which Get operations have been performed are finally deleted. However, in some transactional 10 systems, certain Put_Message and Get_Message operations can be made to take effect immediately without awaiting the final resolution of the transaction - these are referred to as "out-of-syncpoint" Put and Get operations.

As noted previously, a failure may affect a single queue or multiple 15 queues (for example all queues within a specific CF list structure). If multiple queues must be recovered, it is desirable for a single invocation of the recovery process to initiate recovery of all of the affected queues. Improved processing efficiency can be achieved by recreating a set of 20 affected queues and then performing a single recovery unit of work which encompasses restoration of messages and message updates for the whole set of affected queues.

The recovery process has access to and uses whatever log or logs 25 contain information relating to changes to the queue or queues being recovered. In a shared queue environment, it is likely that each queue manager will have maintained its own physically separate log, and each log can comprise a set of files. The recovery process can read all of the logs in parallel, logically constructing a single, merged log. The single merged 30 log (which in general does not exist as a single physical file) contains all of the changes to the queue or queues being recovered, as well as changes to other queues that are unaffected by the failure. The restore process ignores changes to queues which are not required for the current recovery processing.

35 A specific sequence of recovery processing operations are described below in detail, with reference to Figure 3. For ease of reference, the following description of recovery processing describes the example of recovering a single queue.

40 A first step 200 of the method is the identification of a storage failure. In many cases, software using a data repository will be made aware that data has been lost or corrupted by either the hardware (which may be inaccessible, for example) or the operating system or other runtime environment such as a Java Virtual Machine (which may return an error

indication when access is attempted). In the preferred embodiment, the software using a data repository automatically initiates 200 recovery processing when the software becomes aware of a problem. In particular, a queue manager program which is using the failed queue or queues responds to 5 a specific set of error conditions by starting a recovery process which is a component of the queue manager.

In alternative embodiments, the software can be written to present a suitable error notification in response to a failure - prompting human 10 intervention to manually initiate the recovery processing. Additionally, operator action will generally be required to initiate recovery if a storage failure occurs due to accidental or malicious deletion of data.

When initiated in response to identification of a failure, the 15 recovery process accesses secondary storage and retrieves 210 the backup copy of the queue definitions corresponding to the failed queue(s), and uses the retrieved definitions to recreate 210 an empty copy of the queue within primary storage.

20 In the preferred embodiment, the definition of a queue (or other data repository) is held in backup secondary storage separately from the contents of the queue. Backup of the queue definitions as an independent step from backup of a snapshot of the queue contents is beneficial because it facilitates recreation of the queue in an empty state as a separate step 25 before the contents are restored. The queue can be made available for receipt of new messages as soon as it has been recreated in primary storage from its queue definitions.

30 In conventional recovery solutions, a lock is obtained on a newly recreated data repository from the time the repository is recreated until the recovery processing is complete, and locks are perceived to be necessary to prevent duplication of messages. No such lock is required in the present embodiment, and so the data repository (i.e. the queue or database table, but not any updates within the Recovery Unit of Work) is 35 available for use by applications as soon as the data repository is recreated.

Having recreated the queue (in an empty state), a Recovery Unit of Work is then started 230 for restoring messages and message updates to the 40 queue. In addition to the queue definitions required for recreation of a queue in its empty state, the secondary storage contains a backup copy of the queue contents which corresponds to a snapshot of messages on the queue at the time that the backup was taken. The messages within the backup copy are restored 240 to the primary copy of the relevant queue, using a copy

operation together with the step of marking each message to indicate that they are part of the uncommitted recovery unit of work. This marking makes them inaccessible to applications which could otherwise retrieve restored messages from the queue.

5

In the preferred embodiment, the marking of messages is implemented by allocating a unit of work ID and a distinctive primary key to each message, with the value of one byte of the key indicating the state of the message. Queue managers can then interpret the byte value of the primary key to determine whether a message can be retrieved by an application program or not. Any message update within an uncommitted recovery unit of work cannot be accessed by applications at this stage (not until the byte value is changed at commit of the recovery unit of work). This is described in further detail below, under the title 'Distinctive Keys'. The unit of work ID is useful in case the recovery processing is aborted (such as if a queue manager fails part way through recovery processing), since it enables easy deletion of all of the operations performed within the recovery unit of work. IBM Corporation's MQSeries queue manager programs are known to have peer recovery capabilities which enable them to take over queue recovery processing in such circumstances.

As restoration processing proceeds, the recovering queue manager also generates a list of all of the messages for which operations are performed within the recovery unit of work. This list is used later on during commit processing.

Log records, written between the time of the backup copy and the time of the storage failure, are then replayed 250 to provide information about all updates to the queue which have been lost as a result of the failure.

30 Each log record corresponds to a message add operation (such as a Put_Message operation), a message delete operation (such as a destructive Get_Message operation), or a status update (such as a commit or backout). As each log record is replayed, the queue is updated by the corresponding operation and the message is marked with the unit of work ID of the recovery unit of work and by assigning a primary key including a byte value within the 'in-recovery' range of byte values - as described above. This continues until the point in the log records corresponding to the time of the failure.

40 When the restore processing reaches the point in the log records corresponding to the time of the storage failure, the message repository has been restored to the state it was in at the time of the failure - subject to messages added and retrieved independent of the restore process.

At this point, the restore processing is completed by committing 260 the Recovery Unit of Work. A syncpoint is taken to record the consistent state of the queue data and all messages become available to applications. In particular, committing the unit of work includes identifying all 5 relevant updates by referring to the list of messages added, deleted or updated during performance of restore operations for the recovery unit of work and then updating, for each message in the list, the state-indicating byte value within the distinctive primary key to a value representing the new state of the message. Changing the high-order byte value moves the 10 committed messages to a new position in the queue, since the key values are indicative of the desired message retrieval order as well as being indicative of message state.

If the steps of restoring 'old' messages and message updates to the 15 queue fails, the separately performed recreation of the queue should enable the continued use of the queue for 'new' messages while the restore steps of the recovery processing are retried. Thus, the sequence of operations of performing a first recreation step and subsequently reapplying updates by reference to log records not only makes the queue available for new 20 messages at an early stage but also shields the queue recreation and new message processing from any problems affecting the restore processing. The combination of these features can result in significant improvements to the availability of messaging functions as well as avoiding the exceptional processing required in response to 'queue full' conditions.

25

From this point onwards, assuming the recovery was successful, normal message processing operations can continue for all messages on the queue. When a queue manager which is using the restored queue next checks the 30 state-indicating byte value of the message, the new state of the message will determine whether or not it can be retrieved.

An in-syncpoint Get operation within the Recovery Unit of Work differs from a conventional application Get operation in that the new Get 35 operation specifies which message the operation is to retrieve, so as to replay operations from the log in the correct sequence. Conventional Get operations typically retrieve the first available message, but such an approach during recovery processing could result in inconsistencies between the queue at the time of failure and the recovered queue, since a different message may be retrieved by the Get operation during recovery processing 40 than was retrieved by the original Get operation. Therefore, although some applications do not themselves require messages to be processed in the same order as the messages were placed on the queue, nevertheless message updates replayed from log records are applied in a manner which ensures consistency with the sequence of operations performed before the failure.

Suitable techniques for specifying a particular message to be retrieved by a Get_Message operation are already known in the art and so are not described herein in detail. One example implementation is for the Get_Message operation to use the unique key (unique for all messages within 5 a sysplex) which is allocated to each message when the message is added to a shared queue.

Deferral of Restore Operations

10 In the present embodiment of the invention, recovery does not immediately replay in-syncpoint Get and Put operations when processing the log. Instead, as shown in Figure 4, the Get and Put operations are cached 251 until replay of the log enables a determination to be made 252 of the state of the corresponding unit of work. The log is replayed and operations 15 relating to the message queue or queues being recovered are identified. The identified log records are copied to a cache. When the restore processing reaches the point in the log records corresponding to the time of the failure, the cached log records are analyzed 252 to determine the state, at the time of the failure, of each corresponding (original) unit of work.

20 When the determination 252 is performed, one of the following actions is taken:

1. If the unit of work is committed, the Put or Get is performed 256 (as 25 described above) as part of the recovery processing;
2. If the unit of work remains in-doubt at the end of the Recovery Unit of Work, the recovery processing performs the Put or Get but additionally marks the operation as in-doubt 257 and as part of the original unit of work - as required for eventual resolution of the unit of work by the 30 coordinating syncpoint manager; and
3. For all remaining cases (backout, abort, or presume-abort), the cached Get and Put operations are discarded 255.

35 The recovery unit of work is then committed, as described previously.

The recovery processing method described above enables the restore process to run in parallel with use of the newly re-created queue and with 40 efficient recovery processing, without sacrificing assured once-only delivery of messages.

Optimised Handling of Paired Updates

The inventors of the present invention recognised that an
5 in-syncpoint replay of a committed Get operation within the Recovery Unit
of Work is necessarily getting a message Put to the queue within the same
Recovery Unit of Work. The replay may include replay of a Get_Message
operation followed by replay of commit for the original unit of work. The
particular message can be deleted in response to the committed Get_Message
10 operation without waiting for commit of the Recovery Unit of Work at the
end of the restore process. In the present embodiment, Put and Get pairs
within the Recovery Unit of Work are identified 253 and the corresponding
cached log records are deleted 254 from the cache without the need to
update the queue and then delete the update. This feature of the embodiment
15 complements the 'cache-until-resolution' feature mentioned above to avoid
unnecessary processing and to allow the restoring queue manager to reduce
the build-up of messages on the queue. This potentially avoids unnecessary
queue or repository 'full' conditions.

20 Distinctive Keys

It is known within the shared queue support mechanisms of existing
queue managers to use distinctive primary keys to differentiate between
messages in a Coupling Facility (CF) which are in different states.
25 Typically, the states are committed, in-flight and in-doubt. Such uses of
distinctive keys to differentiate between states is described, for example,
in the specifications of commonly-assigned co-pending US Patent Application
No. 09/677,339 and 09/677,341, which are incorporated herein by reference.

30 The present embodiment uses distinctive primary key values for
messages which are in-flight within the Recovery Unit of Work. 'In-flight'
is the state of a transaction before a request is made for commit or
backout (or before a 'prepare to commit' instruction in the case of
two-phase commit). If there is a failure while a transaction is in-flight,
35 the message state is resolved to backout. This is well known as the
"presume abort" approach. 'In-doubt' is a state which applies to two-phase
commit of transactions which involve an external transaction coordinator.
The coordinator issues a 'prepare' request for the transaction to each
resource manager which has an interest. Following completion of the prepare
40 step, the transaction is no longer 'in-flight' but is now said to be
'in-doubt'. Resolution from in-doubt to commit or abort is performed in
response to a subsequent call from the transaction coordinator. Log records
may or may not have been written for Get and Put operations performed by an
in-flight transaction.

The distinctiveness of the primary keys is achieved by using distinct ranges of values for one byte within the primary key. For example, the first byte of the primary key of messages on a Put list (i.e. a list 5 representing the messages which have been Put to the queue) contains a value in the range X'00' through X'09' if the message is committed and a value in the range X'F4' through X'F6' if the message is not committed. The specific allocation of byte values within the state-indicating range of values simply follows the sequence of values within the range to achieve 10 FIFO ordering. Other schemes for allocating distinctive keys are equally possible.

When an application program issues a Get_Message call, the primary key values of messages in the queue are investigated and compared with a 15 list of key ranges to determine the state of the message. The state of a message as reflected by the primary key value determines whether an application can retrieve the message, but the key values also determine the ordering of messages in the queue and so messages for which retrieval is not possible have key values corresponding to the rear end of the queue. 20 This means that simple numerical ordering avoids irretrievable messages whenever retrievable messages are available in the queue.

Using distinctive keys in this way allows a queue manager to selectively access messages in particular states, and permits simple 25 implementation of other functions such as triggering based on the number of committed messages in the queue. By putting special values in the high-order byte of the key, messages which have been added (Put) to the queue but not yet committed are positioned at the rear end of the list, which makes them easy to ignore when a queue manager is performing a 30 Get_Message operation on behalf of an application.

Distinct high-order byte values can be used to differentiate between a number of different states of a message following invocation of a Put_Message operation. For example, a first range of byte values can 35 indicate a message for which a Put has been performed together with the first 'prepare' phase of a two-phase commit, but the Put is not yet committed; whereas a second range of values indicates a message for which the prepare phase of the commit has not yet been performed following a Put.

40 Two new operational states are defined in the present embodiment, with corresponding distinct keys for each operation and message - one byte of each key containing the distinguishing value within a value range which identifies the state. The new states are only applicable to messages placed in the message repository (in this case the CF shared queue) as part of the

restore process. One state corresponds to uncommitted within the original unit of work (the UoW being replayed) and the Recovery Unit of Work, and the second state corresponds to committed within the original unit of work but as yet uncommitted within the Recovery Unit of Work.

5

These new message states and distinctive key values provide the following benefits:

- In-syncpoint Put operations can be replayed by storing the message on the CF with a distinctive key. The distinctive key prevents the message being processed by other processes that perform actions on the queue, and prevents the message from being included in queue depth calculations, among other things. This means that the restore process does not need to cache these Put operations in memory - which considerably reduces the code complexity and the storage occupancy of the restore process.
- Out-of-syncpoint Put operations and commits of in-syncpoint Put operations can be replayed by setting a key value that is distinct from normal out-of-syncpoint activity. This means that the commit of the Recovery Unit of Work can be performed by updating primary key values (replacing a value in a first range of values with a value from a second range corresponding to a different state) without requiring an in-memory or CF administration structure model of the Recovery Unit of Work. Such structures are required in typical alternative implementations.

It will be clear to persons skilled in the art, in the light of this disclosure, that various modifications of the specific embodiments described can achieve the advantages of the present invention and are within the scope of the invention as set out in the accompanying claims.

For example, the above description of preferred embodiments refers to recreating a data repository and restoring data to the repository. It will be clear to persons skilled in the art that some solutions within the scope of the present invention involve restoring all of the data that was in the repository at the time of a failure. Other solutions only require recovery of certain classes of data - such as only recovering persistent messages and excluding non-persistent messages. In the latter, log records may not be written for non-persistent messages such as information-only data broadcasts. For example, a message containing a periodically updated weather forecast or stock price may not need to be recovered if the next update will be available shortly, whereas a message instructing

cancellation of a flight reservation or sale of stocks must be recoverable to enable assured once-only delivery.

Secondly, while the above description noted that processing
5 efficiencies can be achieved by restoring data items to multiple queues within the scope of a single recovery unit of work, alternative implementations will recover each queue within its own separate unit of work. This will decrease the impact of certain types of failure during recovery processing.

10 Thirdly, the above description refers to a specific method for marking messages to make them unavailable for retrieval by application programs until commitment of the Recovery Unit of Work. Other mechanisms for controlling the unavailability of restored messages while avoiding
15 locking the repository for the entire recovery period, are also possible. One such example is setting a unit of work identifier and setting an in-doubt flag for each restored message which is separate from the distinctive primary keys.

20 The above description of a preferred embodiment of the invention uses independently-saved backup copies of a queue's definitions and the queue's contents. Alternative embodiments maintain both the information defining a data repository and the repository's contents at the time of the backup in a single secondary copy. Nevertheless, the recovery processing can retrieve
25 the stored data from secondary (backup) storage and process that data in a sequence to enable a fast recreation of the repository and making it available for new data items, followed by a separate step of restoring the repository's contents.

30 Further embodiments of the invention are applicable to database solutions. In a database table, new rows may be inserted into the table and processed before old rows (which were populated with data prior to the failure) are recovered. During recovery, applications will see the table as containing only the new rows until such time as the recovery is complete.

35 The above description of a preferred embodiment discloses a recovery method which encompasses: (i) rebuilding a data repository in an empty state for fast availability and then handling restore operations as a recovery unit of work; (ii) performing restore operations in dependence on
40 the determined state of the corresponding original unit of work, for efficient restore processing; (iii) optimized handling of paired updates for efficient processing and to avoid build up in the data repository; and (iv) use of distinctive primary keys to indicate specific in-recovery states of data items and updates to data items. While features (i) to (iv)

are complementary, it is not essential to the operation of any one of these features (i) to (iv) for all of the features (i) to (iv) to be implemented together, as will be clear to persons skilled in the art.

CLAIMS

1. A method for recovery from failures affecting a primary copy of a data repository, for use in a data processing system in which updates applied to the data repository during normal forward processing are applied within transactional units of work, the method including the steps of:

5 storing a secondary copy of data representing data items held within the data repository and updates applied to the data repository within said 10 units of work;

15 in response to a failure affecting a primary copy of the data repository, identifying from said secondary copy a set of operations required for restoring said data items and applied updates to a primary copy of the data repository;

20 determining the state, at the time of the failure, of each unit of work corresponding to one or more operations of the identified set of restore operations; and

25 performing restore operations of said identified set for which said performance is consistent with the determined state of the corresponding unit of work, and discarding restore operations of said identified set for which performance is inconsistent with the determined state of the corresponding unit of work.

2. A method according to claim 1, including the steps of:

30 saving to a cache a subset of said secondary copy of data, which subset corresponds to the identified set of operations required for restoring said data items and applied updates;

35 and wherein, subsequent to the step of determining the state of each unit of work, the step of performing restore operations comprises applying restore operations from said cache.

40 3. A method according to claim 2, including the step of deleting from the cache the restore operations for which the corresponding unit of work is determined to be neither committed nor in-doubt, thereby to discard said restore operations for which performance is inconsistent with the determined state of the corresponding unit of work, when performing restore operations.

4. A method according to any one of claims 1 to 3, wherein the step of performing restore operations includes the steps of:

5 performing restore operations for which the corresponding unit of work is determined to be committed; and

10 performing restore operations for which the corresponding unit of work is determined to be in-doubt, and marking the data item to indicate that the unit of work is in-doubt.

15 5. A method according to claim 2, including the step of deleting from the cache any pairs of updates within the set of restore operations, which pair of updates correspond to addition of a data item and retrieval of the same data item and which pair of updates was completed prior to the failure, thereby to discard said pairs of updates when performing restore operations.

20 6. A method according to any one of the preceding claims, wherein storing the secondary copy comprises storing a backup copy of the data repository and storing log records describing updates to the primary copy performed since the backup copy was stored; and wherein the step of identifying said set of operations comprises replaying the log records to identify operations performed on the primary copy of the data repository.

25 7. A method according to any one of the preceding claims, wherein storing the secondary data copy includes maintaining log records that describe operations performed on data items within the data repository, and wherein the step of restoring data to the primary copy of the data repository includes the steps of:

30 replaying the log records of operations performed on data items within the data repository,

35 caching log records relating to operations performed on data items within the data repository within an original unit of work,

determining from the cached log records the state of the original units of work at the time of the failure, and

40 determining, for said operations having cached log records, which operations to perform within the recovery unit of work based on the determined state of the original units of work.

8. A method according to any one of the preceding claims, wherein the data repository is a message repository and the step of restoring data to the primary copy of the data repository comprises performing message add, update and delete operations on the message repository.

5

9. A method according to claim 8, for performance within a messaging communication system, wherein maintaining the secondary data copy includes storing log records to describe updates to the primary copy, and wherein the step of restoring data to the primary copy of the repository includes
10 the steps of caching log records relating to message add, update and delete operations performed under syncpoint control within an original unit of work, determining from the log records the state of the original unit of work at the time of the failure, and determining the operations to perform within the recovery unit of work based on the determined state of the
15 original unit of work as follows:

if the original unit of work is committed, performing the relevant message add, update and delete operations; and

20 if the original unit of work is in-doubt, performing the relevant message add, update and delete operations but marking the operations in-doubt; and

25 if the original unit of work is neither committed nor in-doubt, discarding the cached operations.

10. A data communication system including:

30 data storage for storing a primary copy of a data repository;

secondary data storage for storing a secondary copy of data representing the data repository which secondary data is sufficient to recover the primary copy of the data repository and data held thereon;

35 a recovery component for controlling the operation of the data communication system to recover from a failure affecting the primary copy of the data repository, wherein the recovery component is operable to control the data communication system to perform the steps of:

40 in response to a failure affecting a primary copy of the data repository, identifying from said secondary copy a set of operations required for restoring said data items and applied updates to a primary copy of the data repository;

determining the state, at the time of the failure, of each unit of work corresponding to one or more operations of the identified set of restore operations; and

5 performing restore operations of said identified set for which said performance is consistent with the determined state of the corresponding unit of work, and discarding restore operations of said identified set for which performance is inconsistent with the determined state of the corresponding unit of work.

10

11. A data communication system for transferring messages between a sender and a receiver, the system including data storage for storing a primary copy of a message repository and including secondary data storage, wherein messages are held in the primary copy of the message repository 15 following a message send operation and are retrieved from the primary copy of the message repository for delivery to the receiver, and wherein a secondary copy of the message repository is stored in the secondary data storage and log records are written to record message send and message retrieval events performed within transactional units of work since 20 creation of the secondary copy,

the system including a recovery component adapted to control the data communication system to perform the following steps:

25 in response to a failure affecting a primary copy of the message repository, identifying from said secondary copy a set of operations required for restoring said messages and reapplying message send and retrieval operations to a primary copy of the message repository;

30 determining the state, at the time of the failure, of each unit of work corresponding to one or more operations of the identified set of restore operations; and

35 performing restore operations of said identified set for which said performance is consistent with the determined state of the corresponding unit of work, and discarding restore operations of said identified set for which performance is inconsistent with the determined state of the corresponding unit of work.

40 12. A computer program product comprising program code recorded on a recording medium, for controlling the operation of a data processing apparatus on which the program code executes to perform a method for recovering a data repository from a failure affecting a primary copy of the data repository, for use with a data processing apparatus having a

secondary data storage and having a component for maintaining a secondary copy of data in the secondary data storage which secondary copy is sufficient to recover the primary copy of the data repository and data items held thereon, and wherein updates applied to the data repository are
5 applied within transactional units of work, wherein the method for recovering includes the steps of:

in response to a failure affecting a primary copy of the data repository, identifying from said secondary copy a set of operations
10 required for restoring said data items and applied updates to a primary copy of the data repository;

determining the state, at the time of the failure, of each unit of work corresponding to one or more operations of the identified set of
15 restore operations; and

performing restore operations of said identified set for which said performance is consistent with the determined state of the corresponding unit of work, and discarding restore operations of said identified set for
20 which performance is inconsistent with the determined state of the corresponding unit of work.

13. A recovery component for recovering a data repository from a failure affecting a primary copy of the data repository, for use with a data processing system having primary and secondary data storage and having a component for maintaining a secondary copy of data in the secondary data storage which secondary copy is sufficient to recover the primary copy of the data repository and data items held thereon, wherein updates applied to the data repository are applied within transactional units of work, the
25 recovery component being adapted to perform a method including the steps of:

in response to a failure affecting a primary copy of the data repository, identifying from said secondary copy a set of operations
35 required for restoring said data items and applied updates to a primary copy of the data repository;

determining the state, at the time of the failure, of each unit of work corresponding to one or more operations of the identified set of
40 restore operations; and

performing restore operations of said identified set for which said performance is consistent with the determined state of the corresponding unit of work, and discarding restore operations of said identified set for

which performance is inconsistent with the determined state of the corresponding unit of work.

14. A computer program for controlling the operation of a data processing apparatus on which the computer program executes to perform a method for recovering a data repository from a failure affecting a primary copy of the data repository, for use with a data processing apparatus having a secondary data storage and means for maintaining a secondary copy of data in the secondary data storage which secondary copy is sufficient to recover the primary copy of the data repository and data items held thereon, and wherein updates applied to the data repository are applied within transactional units of work,

the method for recovering including the steps of:

15 in response to a failure affecting a primary copy of the data repository, identifying from said secondary copy a set of operations required for restoring said data items and applied updates to a primary copy of the data repository;

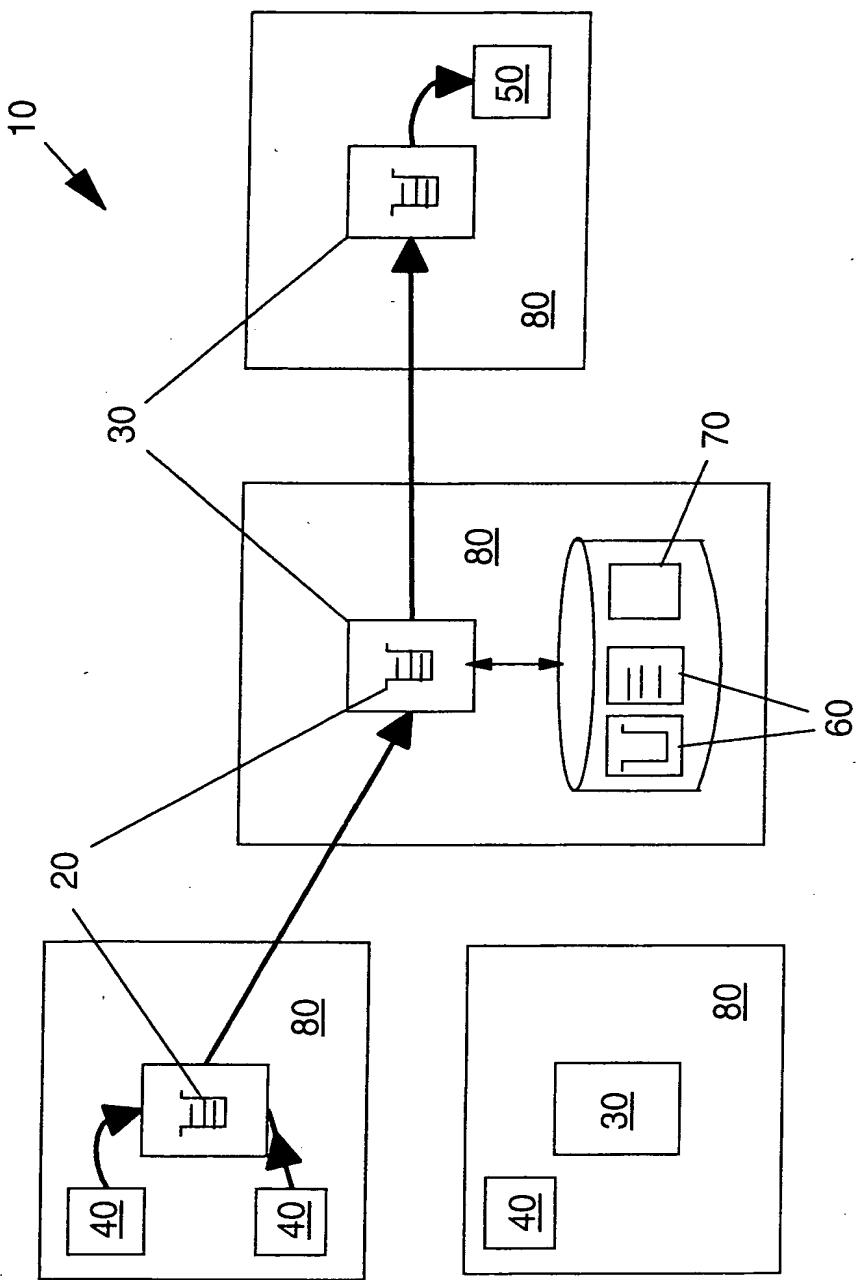
20 determining the state, at the time of the failure, of each unit of work corresponding to one or more operations of the identified set of restore operations; and

25 performing restore operations of said identified set for which said performance is consistent with the determined state of the corresponding unit of work, and discarding restore operations of said identified set for which performance is inconsistent with the determined state of the corresponding unit of work.

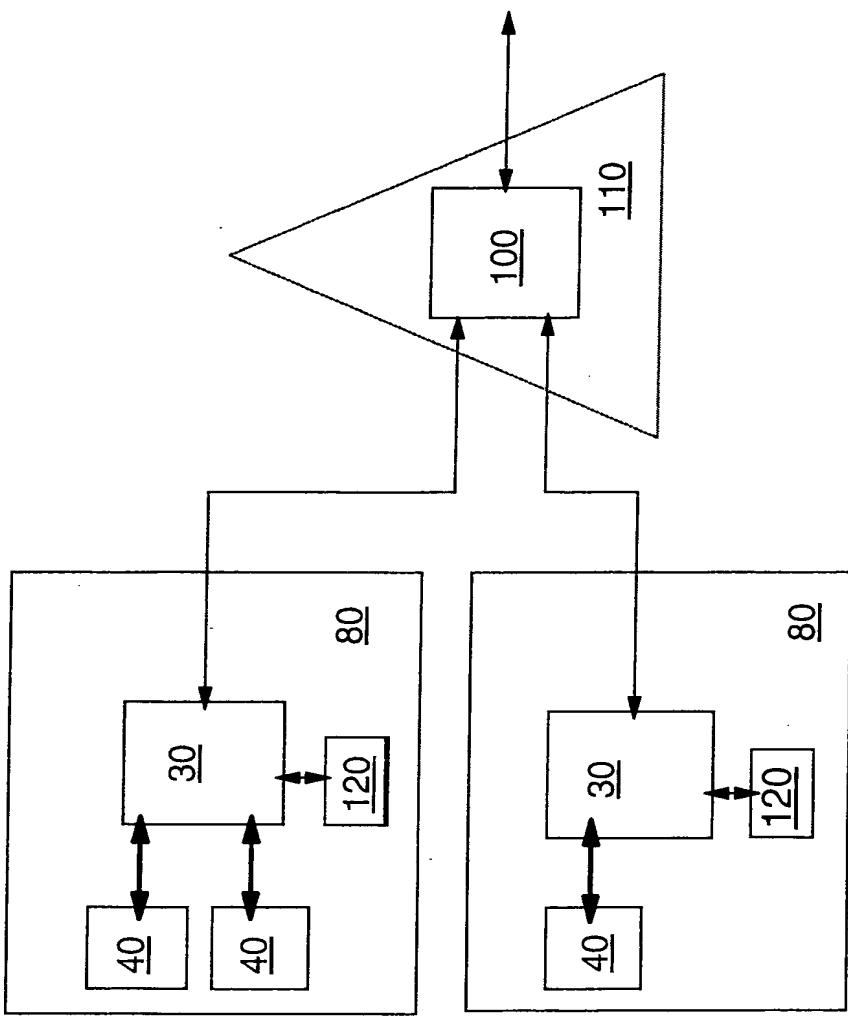
30

ABSTRACT**RECOVERY FROM FAILURES WITHIN DATA PROCESSING SYSTEMS**

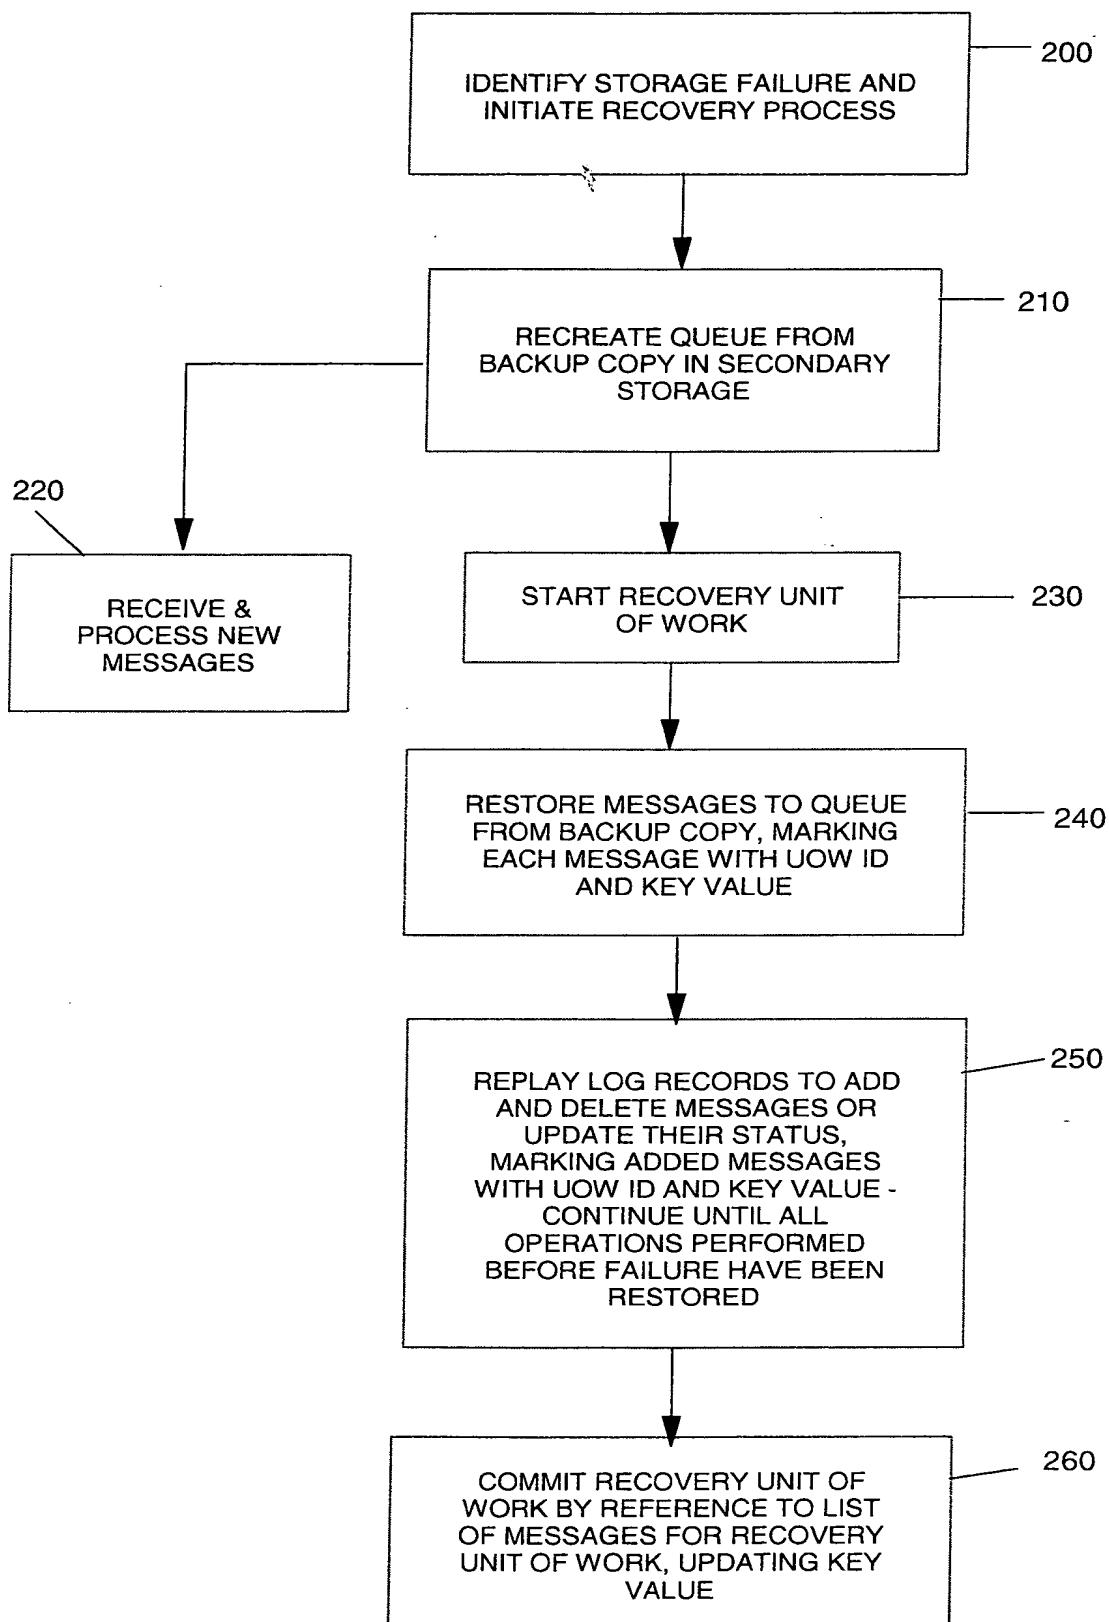
5 Provided are methods, data processing systems, recovery components and computer programs for recovering from failures affecting data repositories. In a data processing system in which updates applied to a data repository are applied within transactional units of work, a secondary copy is stored of data items held within the data repository and updates
10 applied to the data repository within transactional units of work. In response to a failure affecting a primary copy of the data repository, the secondary copy is used to identify a set of operations required for restoring data items and applied updates to the primary copy of the data repository. The set of operations are analyzed to determine the state, at
15 the time of the failure, of each unit of work corresponding to one or more operations of the identified set of restore operations. Restore operations of the identified set are then performed if performance is consistent with the determined state of the corresponding unit of work, but restore operations for which performance is inconsistent with the determined state
20 of the corresponding unit of work are disregarded. The method enables efficiency improvements for recovery processing.

FIGURE 1

1

FIGURE 2

61

FIGURE 3

18

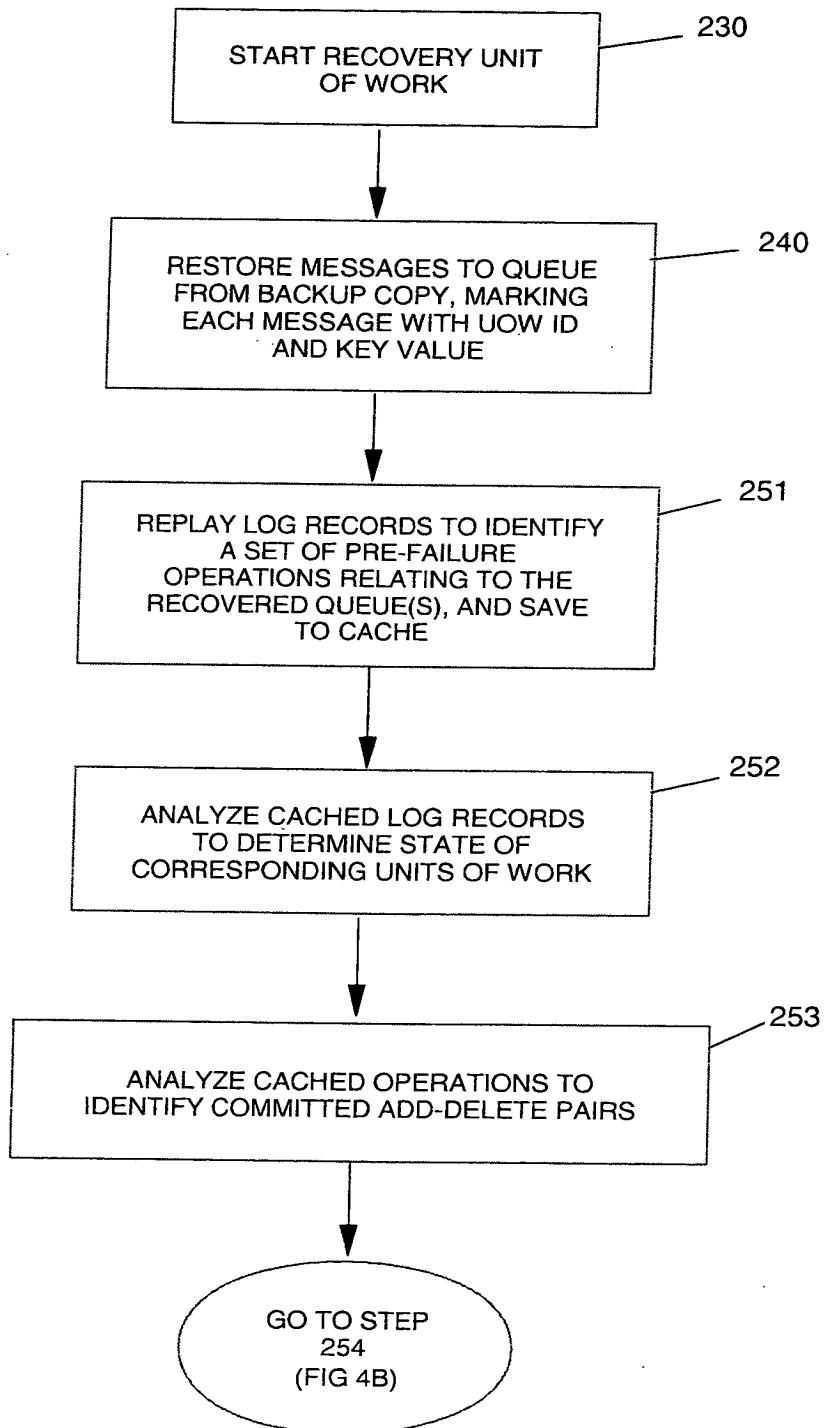
FIGURE 4A



FIGURE 4B